

# PBS Resource Request Examples

When you request PBS resources for your job, you must specify one or more Aitken, Electra, or Pleiades node model (processor) types, so it is helpful to keep this information in mind:

- Usage charges on each HECC node type are based on a common Standard Billing Unit (SBU).
- The total physical memory of each node type varies from 32 GB to 512 GB (with the exception of a few nodes with extra memory, known as bigmem nodes).

The following table provides the SBU rate and memory for each Aitken, Electra, and Pleiades node type:

Processor Type	SBU Rate/Node	Memory/Node
Rome	4.06 (128 cores/node)	512 GB
Cascade Lake	1.64 (40 cores/node)	192 GB
Skylake	1.59 (40 cores/node)	192 GB
Broadwell	1.00 (28 cores/node)	128 GB
Haswell	0.80 (24 cores/node)	128 GB
Ivy Bridge	0.66 (20 cores/node)	64 GB
Sandy Bridge	0.47 (16 cores/node)	32 GB

Note: The amount of memory available to a PBS job is less than the node's total physical memory, because the system kernel can use up to 4 GB of memory in each node.

## Checking Available Nodes Before Requesting Resources

If your application can run on any of the Aitken, Electra, or Pleiades processor types, you might be able to reduce wait time by requesting a type that has more free nodes (nodes currently unoccupied by other running jobs).

Before you submit your job, you can find out the current numbers of used and free nodes for each processor type by running the `node_stats.sh` script, which is accessible by loading the `scicon/cli_tools` modulefile. For example:

```
% module load scicon/cli_tools
% node_stats.sh
Node summary according to PBS:
Nodes Available on Pleiades : 10163    cores: 218436
Nodes Available on Aitken   : 2213     cores: 176272
Nodes Available on Electra  : 3247     cores: 117088
Nodes Down Across NAS      : 1584

Nodes used/free by hardware type:
Intel SandyBridge cores/node:(16) Total: 1668, Used: 1560, Free: 108
Intel IvyBridge    (20) Total: 4705, Used: 4627, Free: 78
Intel Haswell      (24) Total: 1941, Used: 1916, Free: 25
Intel Broadwell    (28) Total: 1790, Used: 1702, Free: 88
Intel Broadwell (Electra) (28) Total: 1066, Used: 1014, Free: 52
Intel Skylake (Electra) (40) Total: 2181, Used: 2165, Free: 16
Intel Cascadelake (Aitken) (40) Total: 1130, Used: 1091, Free: 39
AMD ROME EPYC 7742 (Aitken) (128) Total: 1024, Used: 1024, Free: 0

Nodes currently allocated to the gpu queue:
Sandybridge (Nvidia K80) Total: 59, Used: 32, Free: 27
Skylake (Nvidia V100) Total: 17, Used: 3, Free: 14
Cascadelake (Nvidia V100) Total: 34, Used: 30, Free: 4

Nodes currently allocated to the devel queue:
```

SandyBridge	Total:	72, Used:	0, Free:	72
IvyBridge	Total:	872, Used:	846, Free:	26
Haswell	Total:	91, Used:	52, Free:	39
Broadwell	Total:	400, Used:	381, Free:	19
Electra (Broadwell)	Total:	0, Used:	0, Free:	0
Electra (Skylake)	Total:	100, Used:	100, Free:	0
Aitken (Cascadelake)	Total:	10, Used:	10, Free:	0
Aitken (Rome)	Total:	0, Used:	0, Free:	0
Skylake gpus	Total:	1, Used:	0, Free:	1
Cascadelake gpus	Total:	0, Used:	0, Free:	0

Jobs on Pleiades are:

```
requesting: 14824 SandyBridge, 41922 IvyBridge, 27382 Haswell, 20928 Broadwell,
1559 Electra (B), 6144 Electra (S), 4932 Aitken (C), 0 Aitken (R) nodes
using: 1560 SandyBridge, 4627 IvyBridge, 1916 Haswell, 1702 Broadwell,
1014 Electra (B), 2165 Electra (S), 1091 Aitken (C), 1024 Aitken (R) nodes
```

You can also identify which processor models are in use for each job currently running on HECC systems. Run the following command, and check the **Model** field in the output:

```
% qstat -a -W o=+model
```

## Resource Request Examples

As shown in these examples, use the `model=[san,ivy,has,bro,bro_ele,sky_ele,cas_ait,rom_ait]` attribute to request the node model type(s) for your job.

**Broadwell Nodes:** Both Pleiades and Electra contain Broadwell nodes. If you request only `model=bro` nodes, then by default PBS will run your job on whichever system has available nodes first. If you specifically want your job to only run on Pleiades, then add `-l site=static_broadwell` to your job request. For example:

```
#PBS -l select=10:ncpus=28:mpiprocs=28:model=bro
#PBS -l site=static_broadwell
```

For more information, see [Preparing to Run on Pleiades Broadwell Nodes](#).

The default operating system for all model types is TOSS (Tri-Lab Operating System Stack). For all model types except Aitken Rome, the default is `aoe=toss3`. For the Aitken Rome model type, the default is `aoe=toss4`. You do not need to add `:aoe=toss3` or `:aoe=toss4` to your PBS resource request.

## Example 1: Basic Resource Request

Each of the following sample command lines requests a single node model type for a 128-process job:

```
#PBS -l select=8:ncpus=16:model=san
#to run all 16 cores on each of 8 Sandy Bridge nodes

#PBS -l select=7:ncpus=20:model=ivy
#to run all 20 cores on each of the 7 Ivy Bridge nodes
#(12 cores in the 7th node will go unused)

#PBS -l select=6:ncpus=24:model=has
#to run all 24 cores on each of the 6 Haswell nodes
#(16 cores in the 6th node will go unused)

#PBS -l select=5:ncpus=28:model=bro
#to run all 28 cores on each of the 5 Broadwell nodes
#(12 cores in the 5th node will go unused)

#PBS -l select=4:ncpus=40:model=sky_ele
#to run all 40 cores on each of the 4 Skylake nodes
#(32 cores in the 4th node will go unused)
```

```
#PBS -l select=4:ncpus=40:model=cas_ait
#to run all 40 cores on each of the 4 Cascade Lake nodes
#(32 cores in the 4th node will go unused)

#PBS -l select=1:ncpus=128:model=rom_ait
#to run all 128 cores on 1 Rome node
```

You can specify both the queue type (**-q normal, debug, long, low, or devel**) and the processor type. For example:

```
#PBS -q normal
#PBS -l select=8:ncpus=24:model=has
```

## Example 2: Requesting Different ncpus in a Multi-Node Job

For a multi-node PBS job, the number of CPUs (**ncpus**) used in each node can be different. This feature is useful for jobs that need more memory for some processes, but less for other processes. You can request resources in "chunks" for a job with varying **ncpus** per node.

This example shows a request of two resource chunks. In the first chunk, one CPU in one Haswell node is requested, and in the second chunk, eight CPUs in each of three other Haswell nodes are requested:

```
#PBS -l select=1:ncpus=1:model=has+3:ncpus=8:model=has
```

## Example 3: Requesting Multiple Node Model Types

A PBS job can run across more than one node model type if the nodes are all from Pleiades, all from Electra, or all from Aitken, but not across the systems. This can be useful in two scenarios:

- When you cannot find enough free nodes within one model for your job
- When some of your processes need more memory while others need much less

You can accomplish this by specifying the resources in chunks, with one chunk requesting one node model type and another chunk requesting a different type.

This example shows a request for one Haswell node (which provides ~122 GB of memory per node) and three Ivy Bridge nodes (which provide ~62 GB per node):

```
#PBS -l place=scatter:excl:group=model
#PBS -l select=1:ncpus=24:mpiprocs=24:model=has+3:ncpus=20:mpiprocs=20:model=ivy
```

---

Article ID: 188

Last updated: 08 Sep, 2022

Revision: 76

Running Jobs with PBS -> PBS on Pleiades -> PBS Resource Request Examples

<https://www.nas.nasa.gov/hecc/support/kb/entry/188/>